

**DARPA-BAA-15-04**

**LORELEI**

**Frequently Asked Questions**

As of December 15, 2014

- 126 Q: What is specifically meant by "dictionary" and "lexicon", particularly in the proposers' day slides 24 and 26, as well as in the BAA? Specifically, for TA 1.x performers, will there be a rudimentary pronunciation dictionary for the languages in the LRLPs and/or in the IL to facilitate topic spotting in speech? Or do "dictionary" and "lexicon" have different meanings, more associated with NLP and not with speech processing?
- 126 A: By lexicon we meant an artifact built for use by (text) language technology, whereas by dictionary we were referring to an artifact designed for human use, but in electronic form. The dictionaries might have pronunciation information, but it is not guaranteed, while we do not expect the LRLP lexicons to have pronunciation information.

As of December 11, 2014

- 125 Q: I am on a team with a TA1.4 performer. May I also submit a proposal on one or more of the TA1.1, 1.2 or 1.3 projects separately? Can I propose the same project as I am proposing with my TA1.4 team? Can I propose a different project?
- 125 A: You may propose the same project or a different project separately. Your proposal will not detract from the TA1.4 proposal submitted by the team leader. Of course you will not be awarded a separate award if your project is the same as the project proposed by the TA1.4 team. An individual/organization can only be awarded one contract for any particular amount of a person/group of people's time. One cannot be paid twice for the same work.
- 124 Q: Regarding smaller academic teams applying to LORELEI as TA1.1-3 performers: the BAA states that the best way of demonstrating value to the overall program would be to team with a TA1.4 performer. Does this performer need to be identified ahead of time and listed in the proposal, or is it acceptable to commit to releasing algorithms/tools/resources to the entire community, so that all TA1.4 performers can potentially use them?
- 124 A: The BAA does not state that "teaming" is the best way of showing value, but to "work with a TA1.4 performer to integrate TA1.1, TA1.2, or TA1.3 algorithms, knowledge, resources, or other results into a TA1.4 performer's workflow". Other methods of showing value are acceptable, but must be identified in the proposal. Advance identification of that TA1.4 performer in the proposal isn't necessary. But just "releasing algorithms/tools/resources to the entire community" isn't sufficient, it will be expected that during the course of the program the

TA1.1-1.3 performers will be actively seeking opportunities to "work with a 1.4 performers" or to otherwise show value.

123 Q: What is the security clearance level required for performers on TA2?

123 A: At the time of proposal submission, all proposers to TA2 must have personnel with a Top Secret clearance who are eligible for SCI and eligible for access to facilities to store and process SCI material and hold SCI discussions (see second paragraph on page 13 of the BAA).

122 Q: What is the reasoning behind requiring the TA1.4 performer to also have to respond to TAs 1.1, 1.2, and 1.3?

122 A: It was felt that it would be too risky to decouple 1.4 from 1.1-1.3, and could result in gaps in the program if no one proposed the full suite of TA1 capabilities.

121 Q: On page 43, the BAA says "proposers expecting to use, but not to deliver, commercial open source tools... may be required to indemnify the government against legal liability arising from such use." Under what circumstances would this indemnification likely be required? Does this refer to any third party commercial open source tools? Does it mean specifically software with a viral license, e.g. GPL?

121 A: In accordance with FAR 27.102, contractors are not discouraged from proposing solutions that rely on patented inventions, but should indemnify the Government against liability for the infringement of any U.S. patents. The language in the BAA further extends this principle to open source tools or other materials. Potential proposers are on notice and are responsible for due diligence to ensure the utilization of open source IP in its solutions should not infringe upon the IP of others.

120 Q: Do you see linguistics or ontological inference as distinct areas that could be pursued by TA1 performers? (Linguistic: Knowing that a person must be driving in order to swerve and hit a tree. Ontological: Distinguishing between "drought" vs "guerilla movements" as it pertains to a specific errand and location)?

120 A: Proposers should offer to investigate whatever methods they deem appropriate to address the challenges in the BAA, how they are labeled or categorized is not the key element in how the proposal will be evaluated.

119 Q: Could you further clarify the distinction between TA1.4 and TA2 performers as it pertains to knowledge base fusion – and inference in particular? Specifically are the concepts of "veracity" and "trustworthiness of sources" (conflicts at the KB level) research topics in the scope of a TA1 performer, or should these be left to TA2 in conjunction with stake holders?

119 A: TA2 performers are not expected to need to do language-specific linguistic processing, but they may bring their own general analytics to the effort, including inference or KB fusion, if needed. In general, all single-document linguistic processing is TA1.4 responsibility, and aggregate collection-level non-linguistic processing (including inference and fusion) is TA2 responsibility. Veracity/trustworthiness are not required elements of LORELEI, but could be argued (by a proposer) as being relevant to TA2, in addition to all the stated requirements of TA2.

- 118 Q: Why are evolutionary improvements not funded? A lot of times we start a project, spend time and money to develop it just a little to show it has significant potential and then it is dropped to start a new project as evolutionary developments are not fundable. It seems like wastage of time and money!
- 118 A: DARPA's mission is to develop revolutionary technology. If technology already exists and needs just evolutionary improvement, then the procurement of that technology is outside of DARPA's charter and would be procured by other Government agencies. See the document "Driving Technological Surprise: DARPA's Mission in a Changing World," at <http://www.darpa.mil/WorkArea/DownloadAsset.aspx?id=2147486475>.
- 117 Q: How can multiple performers each do all of 3.1?
- 117 A: All proposals to TA3.1 must cover all of TA3.1, but DARPA may make partial awards. The division could be on the basis of languages, media, resource categories, or something else.
- 116 Q: Is it more appropriate to prepare technologies for a moderate number of languages (e.g. 100) or more general technologies that can scale up to 1000+ languages?
- 116 A: The goal is to develop technologies that apply to a broad range of languages. General tools are more desirable, but DARPA encourages proposals that cover a wide range of solutions.
- 115 Q: In general, if there is some technology that adds value in the incident language, is that alone deemed a success (say NER or summarization), or should each team ensure that results are made available in English?
- 115 A: Technology for rapidly developing a general tool for incident languages could be considered a success, but it is the proposer's responsibility to clearly articulate how it may be useful in the LORELEI context.
- 114 Q: Could the English NER and topic spotting tools that DARPA is willing to provide to TA1.4 teams also be provided to other TA1 performers?
- 114 A: All LORELEI data and tools, with the exception of evaluation data, will be available to all performers.
- 113 Q: If research under TA1 produces linguistic resources (e.g., gazetteers, bilingual lexicons), can those be used by other performers other than through use within a TA1.4 system? Similarly, might performers use resources from other teams (other than through a 1.4 system)?
- 113 A: Yes, to the extent permitted by their licensing restrictions. DARPA desires a minimum of Government Purpose Rights (GPR) for technical data / software developed through DARPA sponsorship
- 112 Q: Many TA1 performers may be able to benefit from some (possibly noisy) gisting or translation capability produced by other performers. Will any access to such translation be available to smaller TA1 performers?
- 112 A: Yes, to the extent permitted by licensing restrictions. DARPA desires a minimum of Government Purpose Rights (GPR) for technical data / software developed through DARPA sponsorship

- 111 Q: Which performers are responsible for deployment, run-time supervision (by an LTDE operator), testing, and updating of LTDE Web services during test exercises conducted within transition partner environments? The BAA states only that the TA2 performer will \*interface\* with TA1-produced LDTE services (pg. 13). Which performers will \*ready, deploy, and monitor\* those services prior to and during the exercises?
- 111 A: Preparation and update of the LTDE Web services for the IL will be the responsibility of the TA1.4 performers. In cases where the test exercise itself is on a restricted network, the TA2 performer will be responsible for installing and maintaining the services on that network.
- 110 Q: How shall TA2 performers distinguish analysis tools that are in scope from language processing capabilities that are out of scope? For example, is it in scope for a TA2 performer to make use of a non-TA1 supplied summarization or entity disambiguation tool?
- 110 A: TA2 performers are not expected to need to do language-specific linguistic processing, but they may bring their own general analytics to the effort. In general, all single-document linguistic processing is TA1.4 responsibility, and aggregate collection-level non-linguistic processing is TA2 responsibility. Collection-level linguistic processing (if not separable into the above two areas of responsibility, e.g., by doing document-level indexing in TA1.4) will need to be allocated to one or the other performer by DARPA.
- 109 Q: What is the anticipated start date?
- 109 A: May 2015.
- 108 Q: What is the program budget? What is the approximate award size for each TA?
- 108 A: No comment.
- 107 Q: What will the interaction with the analyst and the analytic product look like?
- 107 A: Proposers should suggest the form of the interaction. Please refer to the resources listed at the end of the proposers' day slides for information.
- 106 Q: The BAA encourages standalone proposers to TA1.1-1.3, and requires that they explain how they would integrate with a TA1.4 effort. What allowances should TA1.4 proposers make for stand-alone TA1.4 researchers? How many TA1.1-1.4 performers will be expected to be assimilated?
- 106 A: TA1.4 performers are expected to be receptive to new ideas, algorithms, and/or resources from other TA1 performers over the course of the program, if they offer new or improved capability, whether from TA1 performers or from the outside community. There is no fixed requirement as to how many new capabilities each TA1.4 performer is expected to accept. There will be a common API to be defined early in Phase 1. During the course of Phase 1 of LORELEI, TA1.1, TA1.2, and TA1.3 performers are expected to demonstrate value to the overall program goals. The preferred method for demonstrating value is to work with one or more TA1.4 performers to integrate TA1.1, TA1.2, or TA1.3 algorithms, knowledge, resources, or other results into the TA1.4 performer's workflow.
- If a proposer for TA1.1 - 1.3 is seeking a cooperative agreement and chooses not to provide their algorithms, knowledge, resources, or other results to the 1.4 performer(s), the proposal must identify how the team will demonstrate value to the program. DARPA does not anticipate "stand-alone TA1.4 researchers".

105 Q: What is the relative emphasis of MT, extraction, aggregation, and speech technologies?

105 A: Proposers must propose means of solving the situational awareness problem. The relative balance should be as they see fit.

104 Q: Is TA1.4 primarily a software engineering or a research integration activity?

104 A: TA1.4 performer(s) must cover all the requirements of TA1, including research, software engineering, and integration.

103 Q: Will TA2 performers conduct software engineering for TA1?

103 A: No, the TA2 performer will be responsible for producing an analytic tool and interface that only accesses TA1 capabilities through Web services.

102 Q: Does DARPA expect an expert operator for the LDTE on each TA1 team? Will DARPA provide an operator?

102 A: Each TA1.4 performer must provide its own expert operators during the program.

101 Q: Run-time supervision by a general linguistics expert is expected. Does DARPA have requirements for the role of this general linguistics expert?

101 A: Each proposer wishing to use an expert theoretical linguist should propose a role for that person.

100 Q: What Web services should TA2 expect?

100 A: Will be jointly decided by TA1.4, TA2, and DARPA, but will include at least NER, MT, and term/topic tagging.

99 Q: Will DARPA provide mission planners or will TA2 reach out?

99 A: DARPA will provide mission planners for evaluations.

98 Q: What target platforms are anticipated for the deployed technology? A cloud? A desktop? During an emerging incident, will users expect a portal into the runtime system from a handheld?

98 A: The run-time framework is expected to run on a desktop computer (or as a client on a desktop), communicating with the LTDE-produced language tools that are running remotely on a server or distributed framework. Proposers can specify their preferred platform configuration. But it is not expected that there will be any operation on or portal to a handheld.

97 Q: The BAA page 8 and page 11 say that the LDTE is expected to produce Web services to provide language technology capabilities for English newswire. To what extent is it expected that the LDTE will use existing capabilities?

97 A: TA1.4 performers are not expected to produce any new English capabilities. Performers are expected to either find them externally or they be provided to the performers as GFI. The exception is any new analytic that does not currently exist.

96 Q: Will analysts also be trying to assess human terrain features that create danger to responders?

96 A: Yes, this may be one of the included topics in the mission plan.

95 Q: Will TA2 performers be expected to interact with performers other than TA1.4 performers?

95 A: No. TA2 performers are only required to interact with TA1.4 performers.

94 Q: Is there interest in languages that are not traditionally written?

94 A: Yes, if the language is written in some medium. For example, a language without a written tradition but still be used in social media in ad-hoc written forms could be selected.

93 Q: Does the program focus on core technologies or are you also interested in models of human behavior and culture?

93 A: The focus is on human language technologies and their application to rapid response scenarios.

92 Q: How will solo TA1.1-1.3 performers be matched with TA1.4 performers?

92 A: Matching and teaming are up to individual performers. During the program, TA1.4 performers will be encouraged to consider adoption of TA1.1-1.3 products, as well as promising new technologies produced outside the LORELEI program. If performers experience any difficulty in teaming, the Government will assist with assigning TA1.1, TA1.2, and TA1.3 performers to a TA1.4 performer.

91 Q: Can we propose tasks or resources that scale up to 1,000 plus languages?

91 A: The goal is to develop technologies that apply to a broad range of languages. General tools are more desirable.

90 Q: What does "term/topic discovery and classification (for text and speech)" cover?

90 A: DARPA is not specifying the exact nature of this task to allow proposers flexibility to innovate approaches that characterize short utterances/messages in a way that can be mapped to a tractable semantics. Thus, scenario-specific ontology or topic inventory, discovered topics, discovered terms (mapped to English terms or an ontology), manually-specified (and mapped) terms, etc. are all allowable.

89 Q: Will there always be both text and speech input?

89 A: Yes. While the ratio will vary by language and/or scenario, there will always be some text and some speech input.

88 Q: Will the native speaker be available on an on-and-off basis (i.e. not as a single continuous chunk)?

88 A: Yes, if needed, the native speaker will be available on and on-and-off basis.

87 Q: Will there be a different native speaker for TA1.3 or will it be a single one for all performers?

87 A: For the evaluation, the government will provide native speakers. Whether it will be the same or a different individual will depend on availability. The government will also provide limited access to native speakers during research, either the same or a different individual, depending on availability.

- 86 Q: For term/topic discovery and classification, does the Government seek solutions where the topic spotting maps to a scenario-specific ontology, OR discovers and creates topic categories, OR both? How will topic technology be evaluated?
- 86 A: DARPA is not specifying the exact nature of this task to allow proposers flexibility to innovate approaches that characterize short utterances/messages in a way that can be mapped to a tractable semantics. Thus, scenario-specific ontology or topic inventory, discovered topics, discovered terms (mapped to English terms or an ontology), manually-specified (and mapped) terms, etc. are all allowable.
- 85 Q: The BAA states that the funding for TA2 is expected to increase significantly in phases 2 and 3. Is there an expectation that funding for TA 1.1, 1.2, 1.3, and 1.4 would decrease in phases 2 and 3?
- 85 A: No, funding for TA1 work is anticipated to be roughly the same throughout the phases.
- 84 Q: Do we need storage capability for classified documents? Do we need such storage directly onsite, or simply access on our team?
- 84 A: No. Access is the only requirement.
- 83 Q: Do you anticipate all work in this area [TA2?] being classified TS-SCI, or is there an ability to do a wide set of the knowledge elicitation in an unclassified environment for design and development purposes?
- 83 A: We anticipate most TA2 work being unclassified.
- 82 Q: Is cross-document information aggregation in scope for TA2 proposals?
- 82 A: Yes. In general, all single-document linguistic processing is a TA1.4 responsibility, and aggregate collection-level non-linguistic processing is a TA2 responsibility. Collection-level linguistic processing (if not separable into the above two areas of responsibility, e.g., by doing document-level indexing in TA1.4) will need to be allocated to one or the other performer by DARPA.
- 81 Q: For TA3.1, are annotation/translation resources expected to be gold standard quality, or is variable quality expected/acceptable?
- 81 A: Single translation, single annotation, with Quality Control. Not gold standard or reference.
- 80 Q: Is crowdsourcing an acceptable approach to producing annotations/translations for TA3.1?
- 80 A: Yes, crowdsourcing is an acceptable approach for LRLPs.
- 79 Q: For TA3, will pointers to existing open source tools satisfy the NLP tool deliverable (segmenters, POS tagger, morph analyzer etc.), or do tools need to be created specifically for LORELEI?
- 79 A: If open source tools are applicable and available, they may be used to satisfy the deliverable. The TA3.1 performer(s) will be responsible for ensuring that any such tools are appropriate and accessible and delivered as part of the LRLP.
- 78 Q: For TA3.1, is it necessary (or desirable) for data formats, annotation guidelines and general approaches to be consistent across all representative languages?
- 78 A: Yes, it is required.

- 77 Q: For TA3.1, is the data volume for parallel text drawn from the monolingual text documents, or is it in addition to the monolingual text volume?
- 77 A: The parallel text requirement is in addition to the monolingual text requirement.
- 76 Q: Is the ordering of representative languages known/fixed, or can TA3.1 proposals suggest an ordering?
- 76 A: Proposers may suggest an ordering.
- 75 Q: How far in advance will incident languages be known to TA3.2 performer?
- 75 A: Approximately at the beginning of each year.
- 74 Q: Are additional types of annotation (e.g. sentiment annotation) of interest for representative languages?
- 74 A: Proposers may suggest additional annotations, but should be costed as options in the proposal.
- 73 Q: For TA3.2, do all linguistic resources (dictionaries etc.) need to be digital text? If no suitable digital resource is found will scanning/keyboarding be required?
- 73 A: For the resources designated in the BAA as allowable in hardcopy, no scanning/keyboarding will be required.
- 72 Q: For Incident language parallel data, only Incident Language->English data is mentioned. What about Incident Language->another foreign language?
- 72 A: A parallel corpus between the Incident Language and another non-English (high-resource) language may be considered an additional Category II incident language resource.
- 71 Q: Regarding informal monolingual texts and speech, is there an expectation that informal text collected will be (as far as possible) event- or incident-rich, or is generic/broad domain informal text and conversational speech acceptable if a T3.1 performer cannot source specific items (i.e. Yoruba microblogs) should the missing volumes be gathered from other domains or omitted?
- 71 A: Yes, there is an expectation that the informal text collected will be (as far as possible) incident-rich, but not exclusively so. The expectation is that the specific mix of genres will vary opportunistically with language/incident, depending on what is found.
- 70 Q: We assume the T3.1 performer must secure all intellectual property rights to all data required with a world-wide unrestricted redistribution license?
- 70 A: The 3.1 performer is responsible for securing intellectual property rights. The exact nature of the proposed rights or licensing will be considered as part of the proposal evaluation. TA3 performers are responsible for making the resources they deliver available to all LORELEI performers and the US Government, at a minimum.
- 69 Q: The BAA states there are existing tools (POS tagger, Morphological analyzers etc.) available in some languages. Are we to assume that these are considered to be adequate/complete?
- 69 A: No. See the last paragraph on page 55/top of page 56 of the BAA.



- 68 Q: Will the NLP tools which are already developed (e.g. Bengali, Tamil, Hungarian) be made available to gauge the sophistication and quality requirements? Should we assume that these do not need further work/enhancements?
- 68 A: These may or may not be made available ahead of the program start. It should be assumed that they may require some further work/enhancements.
- 67 Q: What guidelines are available for POS tagging, NER tagging? What technical specifications are available for the tools to be delivered? (Morphological Analyzers, Segmenters, etc.)
- 67 A: The guidelines used for the REFLEX language packs will be available along with the technical specifications for the delivered tools. It should be assumed that they may require some further work/enhancements.
- 66 Q: Will technical specifications be made available for any languages, e.g. POS and NE tagsets, complexity of Morphological Analyzers features, etc.?
- 66 A: The specifications used for the REFLEX language packs for POS and NE tagsets, complexity of morphological analyzers' features, etc. will be available. It should be assumed that they may require some further work/enhancements.
- 65 Q: We assume that coref is done on doc level, not sentence level. What are the specifications regarding the documents that will be annotated for coref (i.e. document size, type of text)? Can we assume that a subset of the simple NE tagged text will be used for coref annotation?
- 65 A: Yes, document-level coreference, where the documents are drawn from the NE tagged text, ranging across available genres.
- 64 Q: Regarding topic annotation, is there an expectation that the set of topics will be defined across the entire corpus of speech to be annotated? How many topics are expected to be defined?
- 64 A: Topic labeling will most likely consist of labeling the data into a total of approximately 10-12 topics and sub topics, including relevant and irrelevant topics. The specific topics will be based on what is present in the data.
- 63 Q: Regarding the name transliterator, is there an expectation that this tool will provide only one transliteration for each name, or multiple potential transliterations, depending on possible variations in pronunciation that are commonly reflected in spelling and/or differences in conventions that may be applied in actual texts? Must the chosen transliterations and transliteration system(s) dovetail with those of a given agency or agencies, e.g. the National Geospatial Intelligence Agency?
- 63 A: One (or at most two) transliterations will suffice, in accordance with U.S. Government transliteration standards or standard practices, decided on an IL-by-IL basis.
- 62 Q: Who is the intended audience for the materials developed for TA3 (e.g., subject matter experts, computer programmers for TA1 and 2)?
- 62 A: TA1 and TA2 performers can choose for themselves whether and how to use the TA3-developed materials.
- 61 Q: What is the Interagency Language Roundtable level of the intended audience for the materials developed for TA3?
- 61 A: The materials developed under TA3 are intended for use by other LORELEI performers.

Interagency Language Roundtable levels are not relevant (ILR level 0).

60 Q: What is the encoding of the government-furnished data? Should we assume Unicode/UTF-8?

60 A: Assume UTF-8, but verify.

59 Q: Can we assume all text in the program is represented in Unicode?

59 A: Yes.

58 Q: What subject matter/domain/genres are included in the data? Are they consistent across languages?

58 A: The subject matter/domain/genres will be variable and not consistent across languages except for the "REFLEX English core set," "REFLEX elicitation corpus," and "REFLEX phrase book," and the fact that they will include both relevant and irrelevant data.

57 Q: Has the [REFLEX] data been updated since 2008?

57 A: The REFLEX language packs have not been updated since before 2008. The other government-furnished data has been collected/created more recently, and some has not even been collected/created yet.

56 Q: Is the government-furnished data to be used both for TA3.1 and TA3.2 (e.g., Mandarin, Uzbek)

56 A: With the exception of the evaluation data, any LORELEI performer may use any LORELEI data in any way.

55 Q: What is the difference between the lexicon required for TA3.1 and the parallel and monolingual dictionaries for TA3.2?

55 A: The lexicon in TA3.1 is intended for use by algorithms developed under TA1, so will be in digital form in a standard encoding to be agreed upon. The TA3.2 parallel or monolingual dictionaries refer to found (on the internet or in library) resources for the incident language, and are most likely to be for human use, and may even be in hardcopy.

54 Q: What is the difference between the grammatical sketch required for TA3.1 and the monolingual and parallel grammar books for TA3.2?

54 A: The grammatical sketch is to be created by the TA3.1 performer with the intention that it be used by English-speaking HLT professionals or LTDE operators (humans). The TA3.2 grammar books are found (on the internet or in library) resources, so may range from children's primers, to second-language learner's textbooks, to professional-level grammar descriptions, depending on what is found, and may be in hardcopy.

53 Q: What is the difference between the specialized word lists required for TA3.1 and the gazetteers for TA3.2? What other specialized word lists are required?

53 A: The specialized word lists will contain information about other related topics outside of location and place name information such as terms for geopolitical entities.

- 52 Q: Page 54 of the BAA states that "All numbers are approximate." What is the range of acceptability for number of words?
- 52 A: There is no preset range. The acceptable range will be determined on the basis of availability and may vary across languages.
- 51 Q: Should the government-furnished Information be considered the "gold standard" for the part-of-speech tagged text and the morphologically-analyzed text?
- 51 A: No.
- 50 Q: Are multiple translations required for the parallel text (translated into English and translated from English)?
- 50 A: Multiple translations (gold standard reference translations, or another type of reference) will be required for Category III (evaluation) data only.
- 49 Q: Should the government-furnished Information English core sets (66,000 words) be used for the translations?
- 49 A: The English core set is to be used for the English-to-IL parallel data, thus the TA3.1 performer needs to translate it.
- 48 Q: Will a template be provided for the phrasebooks?
- 48 A: Examples of phrasebooks from the existing (government-provided) languages will be provided soon after the beginning of the program. One example (Uzbek) is already available at <http://www.darpa.mil/opencatalog/BOLT.html>
- 47 Q: Is the "grammatical sketch" different from the "descriptive grammar" cited in the government-furnished information appearing on pages 56-57 of the BAA?
- 47 A: No, those two phrases are meant to refer to the same thing.
- 46 Q: On pages 56-57 of the BAA, under Amharic, Burmese, and Guarani for lexicons the notation appears as "Yes". Please explain.
- 46 A: For the languages listed on pages 56-57, the items marked as "yes" have existing lexicons, but the government does not currently have an exact count of the number of entries.
- 45 Q: For representative languages with mutually unintelligible dialects (like Arabic) is any particular variety being targeted? What variety of Wolof is required?
- 45 A: There is no specific a priori requirement, but it is expected that the dialects/varieties of the LRLP languages will be the most universal or common variety of each language. In the case of Wolof, it will most likely be Dakari.
- 44 Q: How should the word to be counted in the speech sections of Category I be determined without a transcript?
- 44 A: It should be estimated. Historically, word count per hour has been approximately 7,000 words per hour of speech.
- 43 Q: What is the "non-English" language required for the parallel dictionary in Category II?
- 43 A: The choice of non-English language will vary depending on the Incident Language. For example,

in the case of South American languages, Spanish is frequently used in bilingual dictionaries.

42 Q: Who is the intended audience for the monolingual parallel grammar book and monolingual primer book? What is their intended ILR level? How many chapters/pages are anticipated?

42 A: All TA3.1 and TA3.2 data is intended for use by TA1 and TA2 performers. The number of chapters/pages will depend on what is available for each language. The ILR level in the IL of the users is likely 0.

41 Q: Should the three gazetteers listed in Category II be created from the collected or Government-provided corpora or could they be developed from other resources?

41 A: The point of the gazetteers is to provide geographic name information beyond that available from the corpora. These gazetteers need to be found (on the internet or in library) or extracted from other sources.

40 Q: Does "translated 4x" under Category III mean that the text corpora should be translated by four different translators?

40 A: Yes. However, it is possible that the determination will be made during the course of the program to create a different type of reference translation set of comparable cost.

39 Q: What is the extent of the topic labeling required in Category III?

39 A: Topic labeling will be required for all Category III data. It will most likely consist of labeling the data into a total of approximately 10-12 topics and sub topics, including relevant and irrelevant topics. The specific topics will be based on what is present in the data.

38 Q: TA 1.1 - what kinds of resources are expected? How about resources not furnished by the program?

38 A: TA1.1 performers may use any resources they desire

37 Q: Will there be trilingual parallel data?

37 A: There will be a common core corpus across all LRLP languages of about 100K words, which will result in 25-way parallel corpus.

36 Q: In the concept of operations, is resource creation considered part of the 24 hours?

36 A: Yes.

35 Q: Does an IL become an LRLP after it is used?

35 A: An IL does become available after it is used, but funding may not be available to expand it to full LRLP inventory.

34 Q: Compared to resources available for low resource languages, there is a vast array of resources in English, and a significant quantity and diversity of resources already available in Arabic and Chinese. Are performers allowed to use those resources and data derived from them in LORELEI?

34 A: Yes, performers can use available resources in languages other than the language of the evaluation, with the exception of a specific year-long epoch that may be blacked out for each evaluation. During the program, performers must disclose what resources they use, even if those resources are not available to others.

- 33 Q: What redistribution rights for bilingual dictionaries and text corpora (for example) need to be obtained by TA3 performers? Will intra-program sharing of existing ("found") but copyrighted resources be considered fair use (e.g. as derived products)? Otherwise, should proposal budgets include royalty costs for IP rights acquisition for found linguistic resources if not considered "fair use" within the program?
- 33 A: TA3 performers are responsible for making the resources they deliver available to all LORELEI performers and the US Government, at a minimum. The IPR and license terms offered by a proposal will be considered as aspects of proposal evaluation.
- 32 Q: Will the computational tools developed under TA3 be available to all performers?
- 32 A: Yes. They will be released as part of the LRLPs according to the schedule in the BAA.
- 31 Q: How can one gain access to the REFLEX data?
- 31 A: The Uzbek REFLEX data is the core of the data set mentioned in the BAA, and linked in the DARPA Open Catalog at <http://www.darpa.mil/opencatalog/BOLT.html>. The Linguistic Data Consortium at the University of Pennsylvania has published some of the other REFLEX data sets and the rest should be released in the near future.
- 30 Q: During the annual field evaluations, which will need to be directly supported by TA2 performers minimally during each of the checkpoints, will each TA1-TA2 pairing be evaluated in parallel or in sequence?
- 30 A: DARPA anticipates only one TA2 performer, but multiple TA1.4 performers. Each TA2-TA1.4 pairing would be separately evaluated. The format of that process will be determined based on the availability of mission analysts, hardware, etc.
- 29 Q: How many transition exercises will be conducted?
- 29 A: Two or more.
- 28 Q: We understand a need to demonstrate capability within 24-hours of the onset of an event; at the onset (T0) will any data streams (English and Incident Language) be provided from time before the onset of the focal event (e.g., to serve as baseline data)?
- 28 A: Only if it's available in archive form as part of the IL data set.
- 27 Q: Will the field evaluation users have exposure to or training on the tools prior to the evaluation?
- 27 A: Yes.
- 26 Q: Will all evaluations, including the first, have three checkpoints, with all services available at all checkpoints?
- 26 A: The services need to get more complex over the course of the three checkpoints. For the first evaluation, performers will be given more than 24 hours for the 1 day checkpoint and more than 1 week for the second.
- 25 Q: How will machine translation be evaluated in the lab evaluations?
- 25 A: The lab evaluations will be similar to the NIST OpenMT evaluations, but with resource and/or time constraints (still being developed). <http://www.nist.gov/itl/iad/mig/openmt.cfm>

- 24 Q: The BAA states that only program-released data may be used for the evaluations (p7, p16). Is this true for non-IL related resources? For example could language universals discovered by mining all Wikipedia pages for all languages be incorporated into an evaluation system? Could incident specific lexicons mined from high resource languages? Should it be assumed that the only language materials used to develop the LORELEI interface will be those developed under TA3 (e.g., a 10,000 stem lexicon)?
- 24 A: Performers may use other resources for development. Non-IL resources are encouraged for discovering language universals, building incident-related lexicons, and other resource-related efforts. However, for a specific evaluation, a certain epoch (such as a specific year) may be embargoed to avoid excessive exposure of information concerning the specific natural disaster or other incident at the center of an Incident Language scenario. For evaluations, performers must disclose what resources they use, even if those resources are not available to others.
- 23 Q: Do you anticipate restricting use of English-only resources (e.g., large corpora, WordNet, Propbank, etc...)?
- 23 A: The only restriction will be a possible embargo of certain epochs relating to the evaluation topic.
- 22 Q: What material in the IL is available to the performers? Can we use previously-collected material?
- 22 A: Prior to Tzero, performers will not know the language of each evaluation. During the evaluation, performers will use only the incident language material provided for the evaluation by DARPA and the TA3.2 performer.
- 21 Q: Will TA1 performers receive training data of the type described in Category III on page 15 of the BAA for the Incident Language (i.e. NE tagged, Topic tagged)? If so how many words? At which time-points (1 day, 1 week, 1 month)?
- 21 A: For the "surprise" IL evaluations, the TA1 performers will not receive Category III data for that language until after the evaluation. TA1 performers will receive one or more sample Category III data sets early in Phase 1, including for Mandarin Chinese. In operational settings that LORELEI will be eventually supporting, Category III data will be generally unavailable unless created in the LTDE.
- 20 Q: The Category II table on page 14 of the BAA lists 300K parallel IL-English data per incident language. The program resources on page 7 describe the parallel data as "Variable amounts of (archival) parallel corpora in the IL and English or other language". Will there be at least 300K words of IL-English data provided to TA1 performers? At which time-point (1 day, 1 week, 1 month)?
- 20 A: The resources on page 7 are part of a description for an operational CONOPS (concept of operations) once the LORELEI program has succeeded. In operational settings, we will have no control over quantities of data that are available, and the availability of data may be on a streaming basis (in addition to any archived data that may be found). For LORELEI Program ILs (both development and evaluation), Category I and Category II data will be made available at the beginning of the first day, in the quantities mentioned in the BAA, at least for the earliest evaluation. It is possible that for later evaluation(s), Category I and II data may be streamed in to simulate data availability during an actual incident.

- 19 Q: For TA1, what can we assume about what resources will be available when the BAA says that some resources for IL will be available at T zero with additional resources available over the evaluation period?
- 19 A: Category I and Category II data will be made available at the beginning of the first day, in the quantities mentioned in the BAA, at least for the earliest evaluation. It is possible that for later evaluation(s), Category I and II data may be streamed in to simulate data availability during an actual incident.
- 18 Q: If the same performer is selected as a prime for TA3.1 and TA3.2, can a sub on TA3.1 (but not on TA3.2) then also be a prime or sub on a TA1 task?
- 18 A: Yes. Any proposer to this combination of tasks should include in their proposal a mitigation plan.
- 17 Q: If we submit a TA1.n,  $n < 4$ , proposal, and if we are also a member of a team for a TA1.4 proposal (under which we are proposing other content), does one proposal jeopardize the other?
- 17 A: No.
- 16 Q: Can one individual/organization be a sub-contractor for multiple areas/subareas or a sub-contractor for one area/subarea and a prime contractor for another? If so, does the work have to be different, or can it be the same work? Can one individual/organization propose the same work on multiple teams?
- 16 A: Yes, but an individual/organization can only be awarded one contract for any particular amount of a person/group of people's time. One cannot be paid twice for the same work.
- 15 Q: If we submit to TA1.n,  $n < 4$ , do we need to identify a TA1.4 performer who could be willing to integrate our work? Do we need to do this for the proposal (in this case, a proposed TA1.4 performer)? Or do we need to identify a TA1.4 performer after the program has started?
- 15 A: No.
- 14 Q: Can we devise a way to show value that does not involve integration with a TA1.4 performer (which would be detailed in the proposal)?
- 14 A: Yes.
- 13 Q: Will the project sponsor new TS clearances?
- 13 A: Possibly, but proposers to TA2 must have personnel who already have active clearances.
- 12 Q: Is it expected that classified portions of the program will be restricted NOFORN?
- 12 A: This is possible, but has yet to be determined.
- 11 Q: Are TA2 performers expected to bring additional transition partners beyond those already involved in the LORELEI program?
- 11 A: Not required.

- 10 Q: Is it desirable or undesirable for TA1 teams to include SMEs in situational awareness and mission planning as opposed to just focusing on language solutions and deferring other aspects to the TA2 team?
- 10 A: Permissible but not required.
- 9 Q: What is the definition of "small business" for the purpose of LORELEI?
- 9 A: See the Small Business Administration standards for the NAICS code for this procurement (listed on the FedBizOpps page).
- 8 Q: Can foreign participants submit proposals for TA2?
- 8 A: Each proposer should make an individual determination as to their eligibility. TA2 proposers must have appropriately cleared personnel and provide verification that they will be able to access a variety of transition partner sites, some of which may restrict access by foreign nationals.
- 7 Q: Are foreign IRB determinations acceptable?
- 7 A: We do not anticipate human subject research. If you plan to propose human subject research, please include a plan for receiving appropriate United States Department of Defense approval.
- 6 Q: Are proposers to TA1.4 required to have a defense contracting company as a prime, or can they have an academic institution as a prime?
- 6 A: The prime can be any eligible organization.
- 5 Q: Should TA2 proposers be familiar with mission planning tools and processes?
- 5 A: Yes, this would be considered a strength. Please see the references at the end of the proposers' day slides.
- 4 Q: Are there any prohibitions for a prime/sub-contractor participating in both TA1 and TA2?
- 4 A: No.
- 3 Q: Should one abstract be submitted per proposal or per organization?
- 3 A: This is up to the proposer.
- 2 Q: The presentation appeared to refer to the possibility of offering multiple "partial" TA3.1 awards. What would a "partial" TA3.1 award constitute (e.g. a subset of the languages or subset of the data categories)? If there are multiple TA3.1 awards, how would their efforts be coordinated and be synergistic?
- 2 A: All proposals to TA3.1 must cover all of TA3.1, but DARPA may make partial awards. The division could be on the basis of languages, media, resource categories, or something else.
- 1 Q: Since some programs announce the anticipated size of program and/or the anticipated range of awards by Technical Area, could you provide any information regarding LORELEI?
- 1 A: DARPA has not determined specific award amounts. Proposers should propose levels of effort they judge to be sufficient to accomplish their proposed tasks.